# Towards Automated Performance Bug Identification in Python

Sokratis Tsakiltsidis, Andriy Miranskyy, and Elie Mazzawi

arXiv:1607.08506v1 [cs.SE] 28 Jul 2016

**Abstract**—*Context*: Software performance is a critical non-functional requirement, appearing in many fields such as mission critical applications, financial, and real time systems. In this work we focused on early detection of performance bugs; our software under study was a real time system used in the advertisement / marketing domain.

*Goal*: Find a simple and easy to implement solution, predicting performance bugs.

*Method*: We built several models using four machine learning methods, commonly used for defect prediction: C4.5 Decision Trees, Naïve Bayes, Bayesian Networks, and Logistic Regression.

*Results*: Our empirical results show that a C4.5 model, using lines of code changed, file's age and size as explanatory variables, can be used to predict performance bugs (recall = 0.73, accuracy = 0.85, and precision = 0.96). We show that reducing the number of changes delivered on a commit, can decrease the chance of performance bug injection.

*Conclusions*: We believe that our approach can help practitioners to eliminate performance bugs early in the development cycle. Our results are also of interest to theoreticians, establishing a link between functional bugs and (non-functional) performance bugs, and explicitly showing that attributes used for prediction of functional bugs can be used for prediction of performance bugs.

## 1 Introduction

There exists large number of fields where performance of software is critical. For example, mission critical applications, financial and real time systems [1], [2], [3].

Based on statistics from the International Advertising Bureau, mobile ad expenses for 2014 reached $31.9 billion US dollars [4], followed by a 20% growth in 2015. This revenue is entirely managed by real time buying systems.

Real time systems (RTS) are software that subject to time constraints; they have to process information and provide an output within a pre-specified time threshold. A more specific sub-category of RTS is the application under study, which is a real time buying (or bidding) system (RTB), used in the advertisement/marketing domain. The RTB system, also known as Demand Side Platform (DSP), receives an *http* request from the Sell Side Platform (SSP). This request, which is basically an auction being held by the SSP, contain information that is processed by the RTB. The RTB responds back with an advertisement, if some (or all) of the criteria are met. If the RTB system does not respond within a 100 milliseconds time constraint, the opportunity to bid is lost.

Code inefficiencies for this type of software can easily cause reduced throughput, increased latency, and wasted resources. All the above further translate to unreliable software that cannot respond within the desired time threshold (the aforementioned 100ms), leading to financial losses, due to unsuccessful bidding. Also, since performance bugs are not reported as often as functional bugs [2], [5], it is easier for them to accumulate and affect the system for a longer time. In this study (in line with [2], [3], [6]) we will define *performance bug* as a software defect, where after relatively simple changes are applied, the performance of the software is significantly increased without affecting functionality.

**Problem:** Continuous integration practices are gaining momentum and are being used widely, especially by startups [7]. However, lack of time for sufficient performance testing and immaturity of the development team can cause performance bugs to be injected at all times. The problem in such cases, in contrast with regular bugs that affect the system's outcome, is that such types of bugs are not easily diagnosed. Even if they are, developers can easily omit fixing/improving them in favor of delivering new features or other more significant tasks, such as fixing security bugs [5]. We need a way (preferably a simple and tractable one) to detect performance bugs using information that is either readily available to a developer or can be easily obtained. Static analysis and use of attributes extracted from source code repository are promising candidates.

Moreover, we want to understand which factors (i.e., attributes associated with committed code) affect injection of performance bugs, so that the management can take corrective actions. Therefore, our **research question** is:

*How can we detect performance bugs using static code analysis and attributes extracted from source code repository?*

**Approach:** We tackle the research question by performing a case study; our software under study is a RTB system developed over a period of three years. Our goal was to determine if the commit metrics that we gathered, are associated with the probability of performance bugs to be injected. Approximately 2800 file commits (delivered over this three year period) were processed in order to extract a variety of attributes and metrics (analyzed in Section 3.2) for every file that was changed during development.

- *S. Tsakiltsidis and A. Miranskyy are with the Department of Computer Science, Ryerson University, Toronto, Canada. E-mails: stsakiltsdis@ryerson.ca and avm@ryerson.ca*
- *E. Mazzawi is with Addictive Mobility, Toronto, Canada*

These attributes and metrics were categorized into three groups: a) Project attributes, b) Activity attributes, and c) Experience attributes. Then the data have been preprocessed (see Section 3.3 for details). We tried several approaches both in the manner of attributes selected as well as the classification algorithms that were used, in order to achieve the best outcome. We also quantified the contribution of each attribute on the efficiency of each model.

**Contributions:** Our main contributions from this work are: a) explicitly establishing that static code analysis and metrics can be used as predictors of performance bugs, similar to functional bugs, b) understanding the differences between classification algorithms that are commonly used to predict bugs, c) understanding the impact of each attribute on the models' efficiency.

Performance bugs are often caught while timing reference workloads. However this is typically done late in the development cycle. Our approach is complementary; it enables early detection (and thus early elimination) of the performance bugs, since prediction models operate on the source code and do not require code execution.

**Impact:** What we present in this paper is a starting point for practitioners seeking performance optimization and an efficient way of predicting performance bugs. A successful detection will have a significant impact, since in general, performance bugs need much more time to get fixed than non-performance (functional) bugs [5], [6].

**Structure of this work:** We provide a review of related work that has been conducted in this area, and identify the knowledge gap in Section 2. Next, in Section 3, we present the methodology that we followed and the design of our study. Results of the study are given in Section 4. Threats to validity are identified in Section 5. Finally, Section 6 summarizes our thoughts and outcomes, and specifies future research potentials.

## 2 Related Work

There are many studies in the area of defect prediction (see Hall et al. [8] for review of 208 defect prediction models) and quality assurance of software, since this is one of the most significant endeavors during a product's life cycle. Reduced number of bugs assures better quality, thereby improving the product being delivered and allowing for better resource allocation [8], [9].

There are several categories of tools and methods used to detect and predict defects, with the aim of delivering better software [10]. However, static bug finders and statistical models [11] have become the two most prominent categories of defect prediction. The defects can be further categorized to create prediction models in order to distribute the resolution resources more efficiently [12], [13], [14]. At the same time, efficiency, density, and principally the technical debt caused by defects, play critical roles as well [15].

Failures after the release of a product have also momentous effects, especially for large-scale commercial distributions [2]. Even though developers dedicate large amounts of effort and time to testing, they can never be sure that the system is absolutely reliable.

*Performance*

Even though there are many in-depth studies trying to determine the best way of predicting functional bugs, there is not much work done in prediction of bugs causing performance degradation. Below we give an overview of research papers studying performance bugs.

Performance bugs are harder to expose during testing phase, because they do not cause fatal symptoms and do not affect the overall outcome [16]. Furthermore, it is difficult to find the root cause of performance bugs (in comparison with other types of bugs); they also need more time to get resolved [5].

The work that is closest to ours is by Jin et al. [2], which studied a set of 109 real-world performance bugs. They studied the bugs' lifetime from inception to fix, their root causes and introduction mechanisms, in order to create rule-based detectors. We used a comparable approach to create detectors (referred to as "patterns" in our study) with the aim to identify similar performance bugs. They studied software written in Java, C, C++, and JavaScript; ours was written in Python[1]. Moreover, we focused on finding an efficient method to automatically detect performance bugs based on data extracted through the use of patterns (197 real-world performance bugs in our case). Summarizing, Jin et al. [2] focused on analyzing characteristics of performance bugs, while in this study, our focus is on understanding the contribution of each source code attribute, as extracted from the source code repository, to the predicting power of the several machine learning algorithms that we used. Lastly, our general approach and the methodology for creating the prediction model were meant to be easily reproducible in the future. Therefore, our work is complementary.

Automated detection of performance bugs have been implemented in the past [3], [18]. However, the authors used dynamic analysis tools, namely execution traces and historical performance data, to detect slowly executing code. In addition, dynamic analysis tools often require dedicated testing environment to get accurate performance readings. We, on the other hand, are leveraging code attributes extracted from source code repository and automatically detecting performance bugs before the code executes (and reaches production environment). Thus, our work is complementary.

Static analysis tools can also eliminate performance bugs [19]. However, this approach requires knowledge of pattern template, while ours does not (as we omit pattern type info in our models), hence the complementarity of our work.

There exist formal and thorough coding standards for RTS. However, the standards are, typically, language specific (e.g., C [20] and C++ [21]). To the best of our knowledge, no thorough and formal RTS coding standard exists for Python.

Finally, since in our study we relate performance with the real-time nature of the system, one could argue that Python programming language might not be appropriate for such a case. Although this is a problem that falls out of the scope of this study, we understand its importance and validity. Python programming language (with a proper runtime [22]) is used in well-known large-scale RTS, such

1. Programming language does affect performance of the models, as shown by [17].

as eBay, PayPal, and YouTube [22], [23]. The benefits of the language are best summarized by Cuong Do, Software Architect of YouTube: "Python is fast enough for our site and allows us to produce maintainable features in record times, with a minimum of developers" [23].

## 3 METHODOLOGY

In this section we discuss methodology of our case study. The section is structured as follows. We describe our software under study in Section 3.1. The dataset (extracted from the software under study) and its preprocessing are depicted in Sections 3.2 and 3.3, accordingly. Finally, we discuss building of prediction models in Section 3.4.

### 3.1 Software Under Study

The software that we worked on is a RTB application, written in Python behind an asynchronous web framework. The system has been developed from scratch during the last three years. Although the history of the code might sound short, due to the continuous integration practices, the commit count is 2773 and that is only for the production (master) branch. Therefore, we had sufficient amount of information to work with.

The RTB application receives as input *http* requests and responds back within a time constraint (typically, 100 milliseconds), that is set by the SSPs, which the RTB is connected with. The load of the system can grow tremendously, based on various factors, e.g., incoming request rate and current advertisement inventory.

It is the nature of the system that makes it vulnerable to small fluctuations of the response time, and that is where performance becomes critical. Although in common systems availability is measured by subtracting downtime, in our case, we have to append the total of timed-out sessions in order to calculate downtime. During a timeout, the system's functionality is not affected. However, the strict time threshold is not allowing the system to complete a computational cycle successfully.

#### 3.1.1 Infrastructure

Since performance is usually related to workload, we provide some information regarding the infrastructure associated with the software under study. The cloud cluster undertaking the load consists of 22 virtual machines, each containing 32 vCPUs[2] and 60 GB of RAM, each one of them capable of processing more than 3,000 requests per second. The cluster processes, on average, 10 billion transactions per day (going up to 30 billion transactions on a busy day). However, it is worth mentioning that the amount of machines deployed is dynamic and depends on many factors, such as high peak time periods of the year or the number of countries that the current ads are targeting (the larger the number of countries – the higher the load). In such cases more computational power is added in order to accommodate the load.

2. Intel Xeon E5-2680 v2 (Ivy Bridge) Processors.

#### 3.1.2 Spot the Bottlenecks

As a first step of our methodology, we had to identify the bottlenecks of the system. This was achieved with the help of dynamic program analysis tools while reproducing conditions that happen in production environment and, in some cases, executing the analysis on the production environment itself. Although dynamic analysis tools – the well-known *profilers*, have been improved extensively and can provide accurate results, they cannot identify the performance bugs themselves [24]. However, they can help in preliminary analysis and assist in easy verification of the experimental findings.

By identifying the bottlenecks, we were able to find alternative ways that optimize the performance of the system. Most of the times the changes that had to be applied were a few lines of code, which resulted in significant performance improvements. These findings were categorized and classified as *patterns*.

#### 3.1.3 Patterns

We define patterns as coding practices/styles followed (coincidentally or deliberately) by individual developers. These patterns introduced a logic that affected the performance, but not the overall correctness of the algorithm/system. We identified eleven patterns, described in Appendix A, causing performance degradation (PD). Identification was performed by detecting bottlenecks (i.e. code blocks causing PD) using profilers, followed by manual code inspection. All patterns always lead to PD, which we validated on four platforms and two Python versions. We fixed 197 instances/defects, based on these patterns; fixing patterns leads to 20-99% improvement for a given code block; 99% improvement was achieved by fixing pattern described in Appendix A.2.

These eleven unique patterns were generalized and used in scripts that we built, so that we could detect multiple areas of application. This type of static bug finders were used in the past for similar purposes [2], [25]; however, in this study they were mainly used to get an enriched dataset.

The practice of pattern identification was used before by Jin et al. [2] (rule-based detectors) in order to achieve similar results, i.e., detect multiple areas that were "affected" by the same generalized model. However, their patterns were mainly connected to outdated or error prone API's and, in most of the cases, the bugs were product-specific. Contrariwise, the patterns that we identified are more general. Essentially, they can be re-used on any software written in the same programming language (Python) and are more related to the development techniques utilized by the software developers.

*Benchmarking Patterns*

In terms of benchmarking and calculating the performance improvement of each pattern, we thoroughly tested the impact of each one of them. The performance improvements per pattern ranged between 20% and 99%.

Additionally, and in order to validate that a pattern is not endemic to a particular environment (e.g., Operating System, Hardware Platform, or Programming Language

version) we tested them in multiple environments and observed that they lead to similar performance degradation in all environments.

### 3.2 Resulting dataset

Having the patterns as a stepping-stone (and the static bug finder scripts using them), we could move forward and extract a dataset with the help of the version control system. It is worth mentioning that the dataset that we extracted and utilized, contains information on the performance bug injection points. Such information is difficult to extract from a bug tracking system [26], even when dealing with more general types of bugs.

Therefore, our dataset included performance bug injection commits and was further split by file. As a result we had a thorough dataset including the following attributes[3]:

*Project attributes*

i ***Type:*** Indicates either if the entry injected[4] a performance bug into the code or not (not a defect at all).

ii ***Commit ID:*** Unique identifier of the commit. *Total number of commits is 2773.*

iii ***Owner:*** The user that submitted the commit. *13 unique owners under study.*

iv ***File name:*** The file that was modified. Since each commit can contain several files, we split the entries to unique 2-tuples of commit id and file name (henceforth called tuple). In essence, the tuple serves as a unique identifier for a record in our dataset. *338 unique files, 4623 unique tuples (commit id and file name).*

*Activity attributes*

v ***Lines added:*** The number of lines that were added (or modified) for the corresponding file. Note that one modified line will be counted as one line removed and one line added (a-la diff).

vi ***Lines removed:*** The number of lines that were removed (or modified) for the corresponding file.

vii ***File age:*** Age of the file, computed as 'modification date – file creation date'.

viii ***File size:*** Although two separate values were extracted to represent file size, we treat them as one attribute:

- *Source lines of code.* Lines of code without comments or blank lines.
- *Comment lines.* Lines of code representing the developer's comments.

*Experience attributes*

ix ***Maturity/Expertise of the owner:*** The amount of time an owner spent working on the product, computed as 'modification date – first commit date of the owner'.

x ***Time since last commit:*** How much time elapsed since the last commit (in any of the files).

Descriptive statistics for all the numerical attributes is shown in Table 1. Correlation between the numerical attributes is given in Figure 1. As we can see, none of the variables, except[5] 'source lines of code' (SLOC) and 'comment lines', are strongly correlated.

### 3.3 Dataset Preprocessing

#### 3.3.1 Cleansing

We eliminated all files that did not go into production code base, namely: readme files, testing scripts, and help files.

Additionally, we removed a minor 0.2% (9 out of 4623 unique tuples) of "commit id – file name" records related to source code files. These records were outliers, extreme cases. For example, we excluded source files that were moved or removed. To be more specific, the version control system by default identifies directory changes/refactorings as complete removals of the files themselves. Therefore, whenever a file is moved one or more levels up or down in the directory structure, we noticed abnormal numbers of lines added and/or removed. In some of theses cases (and especially in directories including large files) more than 10,000 lines were added or removed on a single commit. The cleansing described above, resulted in a more precise model creation, mainly because of the removal of entries that will not be seen on future commits [31].

#### 3.3.2 Re-balancing

After the initial cleansing we had our dataset in its final form. However, and as a result of the large number of commits that we have under study, it was expected that the non-defective entries would prevail over the defect set. In fact, we had 95.5% (4426) of non-defective entries and 4.5% (197) of defective ones. This type of imbalance is common, especially when dealing with real world problems [32]. Therefore, in order to avoid biased results towards the majority class, we had to choose a re-balancing technique.

There exists several re-balancing methods, since this ubiquitous problem of imbalanced data is faced quite often. To overcome this obstacle and achieve valid results, researchers use over and under-sampling techniques, either in the minority or majority classes. While reviewing these methods and techniques that have been used towards this goal, we decided to use random over-sampling for our training dataset. Against the numerous (and in some cases combined and complex) methodologies, as studied by Batista et al. [33], random over-sampling provides competitive results.

#### 3.3.3 Overcoming Premature Over-fitting

Given high imbalance of our data, we expected that models trained on rebalanced datasets would be prone to overfitting (as seen by [32], [34]). This was confirmed during

3. This is not an exhaustive list of attributes that can be used to predict injection of a defect [27], [28], [29]. However, we focused on the attributes that can be easily extracted by practitioners to simplify institutionalization of the prediction models.

4. We manually traced injection point for each performance bug using history information obtained from the source code management system.

5. Which is typical for a software product.

TABLE 1
Descriptive statistics of numeric attributes.

| Attribute | Range | Mean | Median |
|---|---|---|---|
| Lines added | 0 - 642 | 19.74 | 3 |
| Lines removed | 0 - 995 | 13.21 | 2 |
| File age (days) | 0 - 803 | 139.2 | 59.97 |
| SLOC | 0 - 1109 | 254.2 | 150 |
| Comment lines | 0 - 419 | 71.07 | 30 |
| Maturity/Expertise (days) | 0 - 997 | 416.9 | 435 |
| Time since last commit (hours) | 0 - 266.7 ($\approx$ 11 days) | 10.71 | 1.26 |

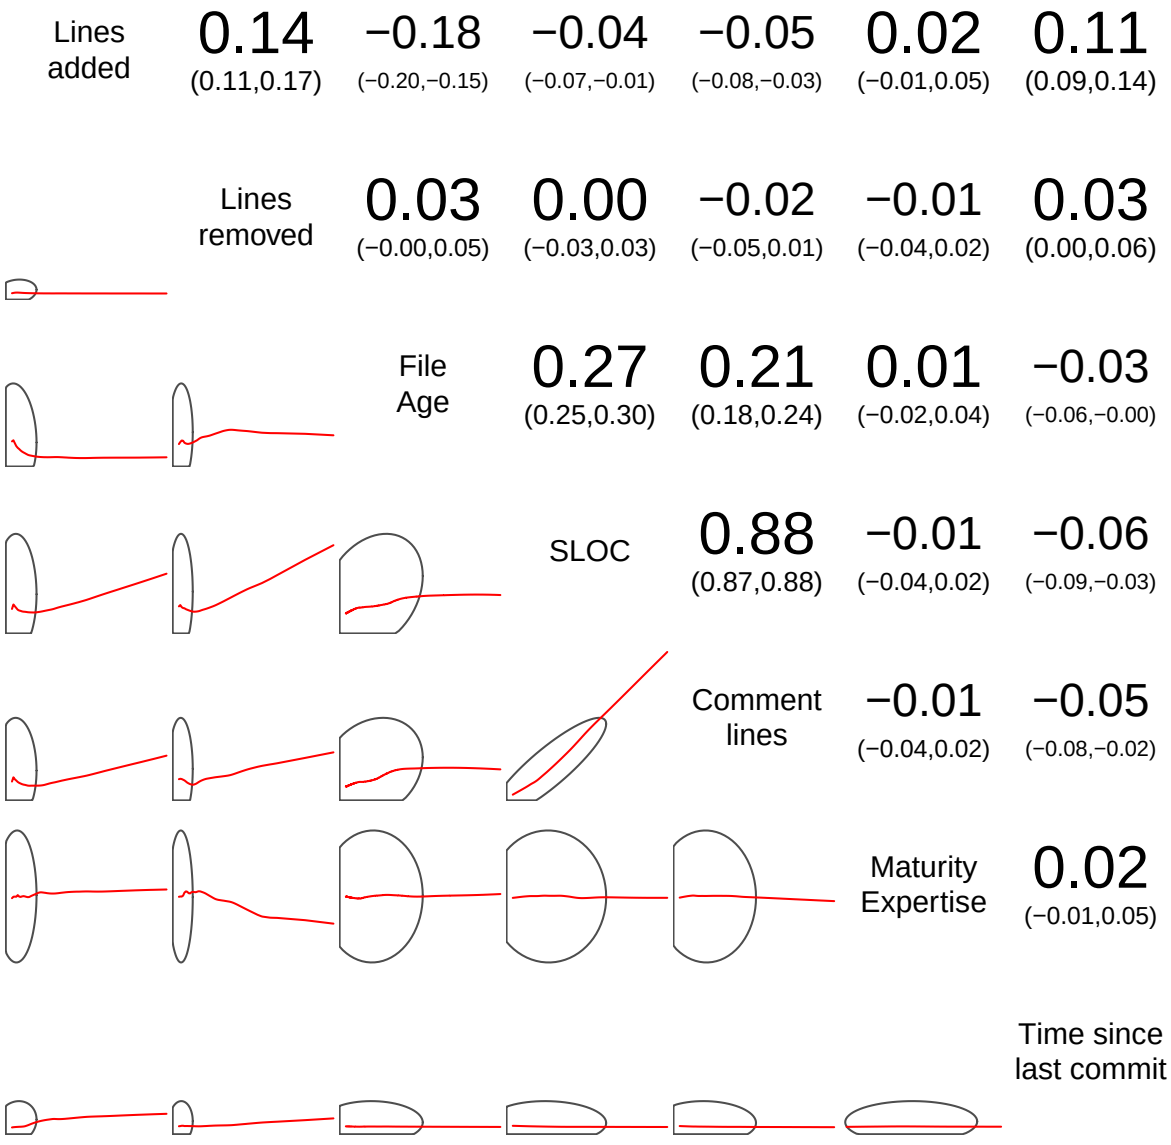


Fig. 1. Correlogram visualizing correlation matrix [30] for numeric attributes. Diagonal shows attribute name. Lower triangle region shows confidence ellipse and smoothed line. Upper triangle region shows Pearson's correlation coefficient $\rho$ and, in brackets, confidence interval for $\rho$.

preliminary model building (discussed in Section 3.4): we observed over-fitting of the models, no matter the classification method.

As described by Tan et al. [35], over-fitting may occur even when one utilizes cross-validation, especially when dealing with extremely imbalanced datasets. Even after applying the re-balancing methods (described in Section 3.3.2), the repetition of the entries might still result in false high precision [35].

To overcome this issue, we used a combination of techniques, similar to the one that was introduced by Mease et al. [34] and called "JOUS-Boost". They added independently and identically distributed noise to the minority class for which random oversampling created replicates. This approach was shown to be superior in dealing with the oversampling problem [32].

In our case, we decided to generate a "jittered" train set, after we split our data to 70% train - 30% validation. We added noise to the minority entries of the 70% part, by randomly increasing or decreasing the values by 0 to 2%. Formally, a variable $x$ becomes $x \times U(0.98, 1.02)$, where $U(0.98, 1.02)$ is a random value drawn from the uniform distribution with minimum value of $0.98$ and maximum value of $1.02$. This transformation could only apply to numerical attributes (listed in Table 1). We do not add noise to the validation set, to mimic the actual defect prediction process; when a developer passes information about a new commit to the prediction 'oracle' to predict if this commit contains a performance bug or not. Results of this experiment are discussed in the next section.

### 3.4 Building the Models

#### 3.4.1 Classification Algorithms

After re-balancing the dataset, it is time to move forward and start building the desired models. The first question that arises is: which classification algorithms should be used for building the models? Based on the literature and existing empirical work [36] we decided to experiment with four popular machine learning methods (using Weka [37] and R [38]). The algorithms chosen are: a) C4.5 decision tree (Weka's J48 class [37]), b) Naïve Bayes, c) Bayesian networks, and d) Logistic regression.

#### 3.4.2 Training and Testing the Models

We trained and tested the model on the 70% of the data using 10-fold cross validation; we then validated the model on the remaining 30% of the data (see Section 3.3.3 for description of data subsets). We used this approach to generate more precise efficiency metrics for the models. The details of the efficiency metrics — namely, accuracy, precision, and recall — are given below.

**Accuracy:** accuracy reflects the percentage of correctly classified performance bugs to the total number of performance bugs. It is computed as follows [39]:

$$ACC = \frac{(TP + TN)}{(TP + TN + FN + FP)}, \quad (1)$$

where $TP$ is the number of true positive results, $FP$ is the number of false positive results, $TN$ is is the number of true negative results, and $FN$ is the number of false negative results.

**Precision:** precision, or Positive Predictive Value is the fraction of retrieved instances that are relevant, and it is calculated as [39]:

$$PPV = \frac{TP}{(TP + FP)}. \quad (2)$$

**Recall :** recall, also known as sensitivity or True Positive Rate, is the fraction of relevant instances that are retrieved, and it is calculated as [39]:

$$TPR = \frac{TP}{(TP + FN)}. \quad (3)$$

All three metrics range between 0 and 1; the closer the value to 1 – the better.

## 4 Results

The baseline of the attributes used for building each model, is the *lines added* to each file on a single commit. Moving forward, new attributes are added (cumulatively) and their impact on the efficiency can be determined by the change of the metrics.

The performance of the models is given in Table 2 and Figures 2-5. Bayesian Nets model yields the best results in the case of a single explanatory variable (*lines added*) with $TPR = 0.68$, $ACC = 0.82$, and $PPV = 0.95$. However, as more explanatory variables are added to the model, C4.5 becomes the top performer. If we exclude the models that use *file name* as the explanatory variable, the best model is based on C4.5 that is using *lines added*, *lines removed*, *age*, and *size* attributes as explanatory variable: $TPR = 0.73$, $ACC = 0.85$, and $PPV = 0.96$. Naïve Bayes and Logistic Regression models, in most cases, take second and third place in the performance competition (surpassing Bayesian Nets).

We will return to the models using *file name* at the end of this section. Let us now analyse which factors affect injection of performance bugs, by analyzing the Logistic Regression (LR) models (because they are the most tractable and reasonably powerful) so that the management can take corrective actions. All the attributes in the LR models are statistically significant: $p$-values $< 0.01$. Recall and Accuracy of the LR models increase as illustrated by positive slopes in Figures 2 and 3, respectively. Precision of the LR models decreases with the addition of the attributes, with the exception of the *file name* attribute (which we will discuss at the end of this section in details), as shown in Figure 4.

As mentioned above, *lines added* attribute has the highest predictive power. Peculiarly, addition of *lines removed* to the model does not improve its predictive power significantly. This can be explained by the fact that bugs are mostly injected in new code rather than in modified[6] or removed one. The sign of *lines added* regression parameter is positive. This implies that in order to reduce the number of injected performance bugs, the management has to enforce reduction of the amount of changes delivered on a single code commit. This aligns with the findings that the data on the number of lines

6. A modified line is counted as one line removed and one line added.

TABLE 2
Experiment metrics: JOUS-Boost. The models are cumulative: the first model describes the model using one explanatory variable: lines added; the second model – two variables: lines added and lines removed; the third one – three variables: lines added, lines removed, and age; and so on. Size consists of two separate attributes, as described in Section 3.2.

| Row # | Variable | Metric | C4.5 | Naïve Bayes | Bayesian Nets | Logistic Regression |
|---|---|---|---|---|---|---|
| 1 | +Lines added *jittered* | TPR: | 0.65 | 0.52 | 0.68 | 0.67 |
| | | ACC: | 0.81 | 0.74 | 0.82 | 0.79 |
| | | PPV: | 0.95 | 0.93 | 0.95 | 0.89 |
| 2 | +Lines removed *jittered* | TPR: | 0.67 | 0.52 | 0.67 | 0.67 |
| | | ACC: | 0.82 | 0.73 | 0.82 | 0.79 |
| | | PPV: | 0.95 | 0.91 | 0.95 | 0.89 |
| 3 | +Age *jittered* | TPR: | 0.70 | 0.52 | 0.67 | 0.68 |
| | | ACC: | 0.84 | 0.73 | 0.82 | 0.80 |
| | | PPV: | 0.96 | 0.91 | 0.95 | 0.89 |
| 4 | +Size (SLOC & Comments) *both jittered* | TPR: | 0.73 | 0.83 | 0.58 | 0.70 |
| | | ACC: | 0.85 | 0.73 | 0.79 | 0.80 |
| | | PPV: | 0.96 | 0.69 | 0.98 | 0.88 |
| 5 | +Experience *jittered* | TPR: | 0.73 | 0.83 | 0.58 | 0.70 |
| | | ACC: | 0.85 | 0.74 | 0.79 | 0.80 |
| | | PPV: | 0.96 | 0.70 | 0.99 | 0.87 |
| 6 | +Time since last commit *jittered* | TPR: | 0.72 | 0.88 | 0.52 | 0.70 |
| | | ACC: | 0.84 | 0.72 | 0.76 | 0.80 |
| | | PPV: | 0.96 | 0.67 | 0.99 | 0.88 |
| 7 | +Owner | TPR: | 0.68 | 0.90 | 0.52 | 0.77 |
| | | ACC: | 0.83 | 0.73 | 0.76 | 0.83 |
| | | PPV: | 0.97 | 0.68 | 0.99 | 0.87 |
| 8 | +File name | TPR: | 0.83 | 0.88 | 0.57 | 0.90 |
| | | ACC: | 0.90 | 0.85 | 0.78 | 0.91 |
| | | PPV: | 0.96 | 0.84 | 0.99 | 0.92 |

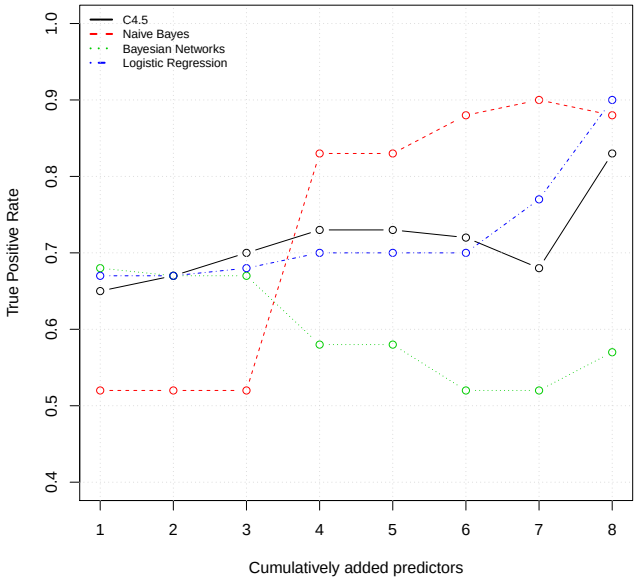


Fig. 2. Models' efficiency in term of Recall (True Positive Rate) for cumulatively added predictor variables. The labels of the $x$-axis correspond to the values in the 'Row #' column of Table 2.

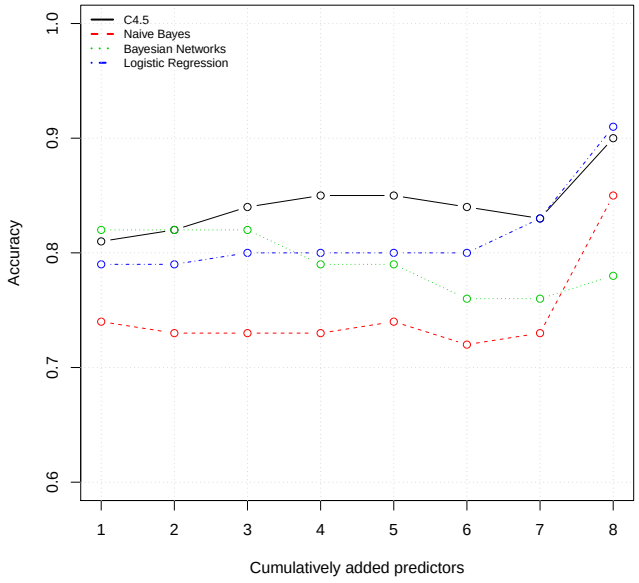


Fig. 3. Models' efficiency in term of Accuracy for cumulatively added predictor variables. The labels of the $x$-axis correspond to the values in the 'Row #' column of Table 2.

of code changed provides most of the information needed to predict functional bugs [40]. In essence, the more lines you change[7], the higher the probability that you will inject a bug. Our findings show that this statement is true not only for general functional bugs [40], but also for performance bugs.

The attributes *age* and *size* improve performance of the models, but not significantly: e.g., compare $TPR = 0.67$ with $TPR = 0.70$. The sign of *age* regression parameter is

7. In our case, we can narrow the type of change to code addition.

positive. From the management perspective, this suggests that a performance bug has higher probability of injections, as the software gets older, which is consistent with Lehman's laws of software evolution [41]. The signs of *size* attributes *SLOC* and *comments* are negative, suggesting that larger files may have lower chances of injecting performance defect. However, given the low improvement in predictive power, the model may be capturing correlation rather than causality.

Merit of *experience*, *time since last commit*, and *owner* is also arguable. They either do not enhance the predictive power of the model significantly or, in some cases, reduce it, as shown in Table 2 and Figures 2-4. From the management perspective, this suggests that all developers (independent of their experience with the code base) have injected performance bugs in the past.

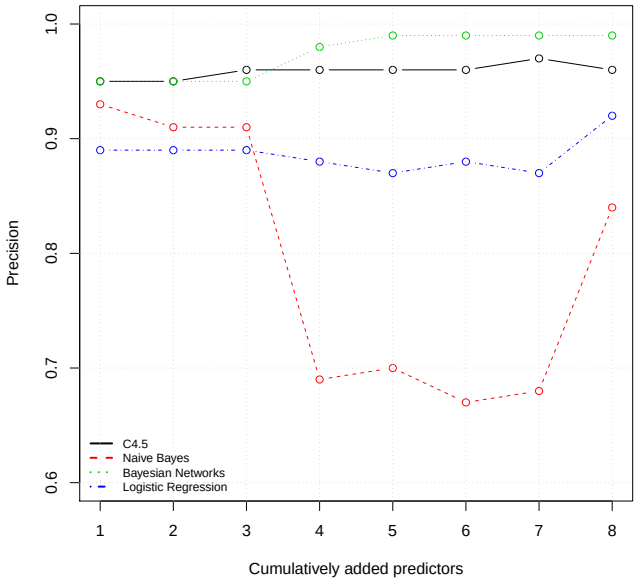


Fig. 4. Models' efficiency in term of Precision for cumulatively added predictor variables. The labels of the $x$-axis correspond to the values in the 'Row #' column of Table 2.

Inclusion of the *file name* improves the performance: the best model is based on Logistic Regression with $TPR = 0.90$, $ACC = 0.91$, and $PPV = 0.92$. This is not surprising, given that our performance bugs reside in 9% (32 out of 338) of the files. It implies that some files are more prone to performance bugs than others (as seen in other projects [28]). The management can focus quality assurance efforts on these 9% of the files, e.g., making sure that code commits to these files are peer reviewed before merging them into the production code base.

However, if we reuse our model on a new project, inclusion of the file name attribute into the model may not be practical, since, at the beginning, information about files and new bugs will be limited [42]. In this case we can resort to the general C4.5 based model mentioned above: it has lower recall and accuracy, but higher precision.

In the final model selection, comparison of True Positive and False Positive rates, shown in Figure 5, might be of help. In this graph[8], we can measure the performance based on the True Positive rate, taking under consideration the drawback of the False Positive rate identification of the models. Since each case and application should be treated differently, we cannot explicitly propose a single model as the best performing one. We will discuss our final selection in Section 6.

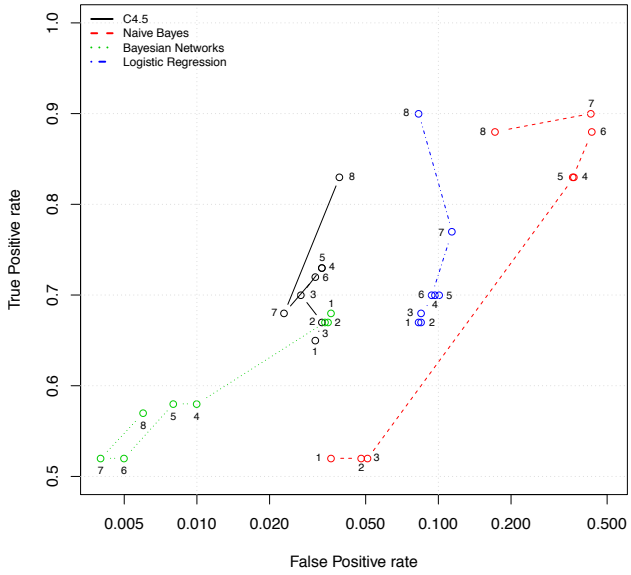


Fig. 5. Models' efficiency in term of True Positive Rate vs. False Positive Rate. The points on the graph represent models with different number of cumulatively added predictor variables. Points' labels correspond to the numbers given in the 'Row #' column of Table 2.

## 5 Threats to Validity

A number of tests are used to determine the quality of case studies. In this section we discuss four core tests: internal, construct, statistical, and external validity [44].

**Internal:** One of the authors of the study is affiliated with the company. In order to avoid researcher bias, we derived and followed strict automated processes for data extraction and processing. The results of extraction and processing were cross-validated by a researcher not affiliated with the company. One of the most critical internal parts of our study is the patterns identification and validation, as explained in Section 3.1.3. Validating, benchmarking and leveraging the patterns was one of the most important parts of our study. We tested the patterns in multiple environments (by varying Operating System, Hardware Platform, and Programming Language version) and observed that the patterns lead to software slowdown in all environments.

Additionally and as already discussed in this paper, the real-time constraints that are externally set for the system under study, as well as the programming language that is used, might lead to bad interpretation of its real-time

8. This graph similar in nature to the receiver operating characteristic (ROC) curve [43].

dynamic. However, this study begun and was applied on an existing, already developed and fully functioning software.

**Construct:** We construct our dataset based on the data collected from the source code management system. The system does not capture all events and activities happening in the organization. However, based on our results, even this incomplete set of data has strong predictive power.

**Statistical:** To prevent over-fitting of the models, we utilize 10-fold cross validation for model creation and the JOUS-Boost [32] approach.

**External:** Generalizing our findings from a single-case design study to all situations is obviously not possible. However, this design is based on the rationale of the critical case [44] of a Real Time System. Our results should be transferable to other researchers with well-designed and controlled experiments. In addition, the analysis can be replicated on open source software.

## 6 Conclusion

In this work we focused on detecting performance bugs that are important in multiple fields, such as mission critical applications, financial, and real time systems. Our software under study is a real time system used in the advertisement / marketing domain.

Our research question was: *How can we detect performance bugs using static code analysis and attributes extracted from source code repository?* To answer the question, we described and followed a methodological approach that resulted in application of four different prediction algorithms, namely: C4.5 Decision Trees, Naïve Bayes, Bayesian Networks, and Logistic Regression.

The best model (not taking into account information about file names[9]) is obtained with C4.5, using lines of code added, lines of code removed, file age, and file size as explanatory variables ($TPR = 0.73$, $ACC = 0.85$, and $PPV = 0.96$). The best model, which has file name available to it, is based on Logistic Regression with all the attributes included: $TPR = 0.90$, $ACC = 0.91$, and $PPV = 0.92$.

Furthermore, the analysis of the variables, based on Logistic Regression model, can aid management in setting corrective actions. In particular, injection of performance bugs can be reduced by decreasing the amount of changes delivered on a single commit[10], as well as by focusing quality assurance resources on a small subset of error-prone files (9% of the total number of files).

We believe that our methodological approach is of interest to practitioners, because it provides them with a simple and tractable model, using easily extractable code attributes for predicting performance bugs on a new code commit. In our models we omit pattern type information to show that the data, extracted automatically from version control system, is sufficient to predict performance bugs. This suggests that performance bugs mapped to non-functional requirements can be detected using the same attributes (particularly "lines added") as functional bugs. This simplifies adoption of the models, since complexity and lack of human and/or hardware resources are often discouraging reasons. It also speeds up detection of a performance bug, as the model is applied to new code before it was committed to public code branch. If the model detects the performance bug, developer decides whether to review the code deeply or proceed as-is.

The resulting factors affecting projects may vary from project to project; however, practitioners may replicate model creation process, described in this study, tailoring it to their needs. Moreover, we describe how to process an extremely imbalanced dataset (where the number of non-defective records is $> 20$ times larger than the number of defective ones), which can help practitioners facing similar issues.

The pattern identification that we conducted and briefly described in this paper, can also be considered as a guide for performance oriented best practices for the Python programming language. Additionally, these patterns can also be reused as static code analysis tools for identification of performance related fixes. They can also be considered as part of enhancement and development guidelines (e.g., Python's enhancement proposals [47]) for Python language. Based on the methodology that we described, researchers can replicate similar pattern identification for other programming languages.

Our results are also of interest to theoreticians, since this work establishes a link between functional bugs and (non-functional) performance bugs, explicitly showing that attributes used for prediction of functional bugs can be used for prediction of performance bugs as well.

Going forward, we would like to extend our work to include other projects and company data.

## Appendix A
## Patterns

### A.1 Pattern 1

An example of a Coding pattern is as follows. Python provides multiple methods to concatenate strings. For example,

```
foo =  'a_string'
bar = 'abc' + 'def' + str1
```

is much slower than

```
foo = 'a_string'
bar = '{0}{1}{2}'.format('abc', 'def',
    str1)
```

### A.2 Pattern 2

An example of a Design pattern is as follows. Python provides a default logging mechanism, which, in some cases, cannot avoid expensive function calls and/or calculations, even if these are finally not necessary/logged. A simple workaround is by doing a cheap function call before executing the actual log function. Essentially, the "cheating" function is just making sure that the current line will be logged before the execution of the expensive call. For example,

9. A novel project may not have sufficient amount of file-name-related data [42].

10. This will lead to increase of the number of commits. However (based on the authors' industrial experience as well as the literature [45], [46]) individual commits will become easier to review, leading to improved defect removal.

```
logger.setLevel(logging.CRITICAL)
logger.debug(doSomethingExpensive())
```

the above example will spend all the time necessary to call the `doSomethingExpensive()` but finally won't use any of this information due to lower current logging level. However,

```
logger.setLevel(logging.CRITICAL)
if logger.isEnabledFor(logging.DEBUG):
    logger.debug(doSomethingExpensive())
```

preliminary check of the log level can prevent unnecessary calculations. In case of the if-block returning `FALSE`, the expense of calling `doSomethingExpensive()` is saved. We made sure that the combination of an if-block and the call of the `isEnabledFor(<level>)` is always cheaper or, at the very least, the same as the default logging function. It might be worth mentioning that the application of the above pattern, resulted in a very significant ($> 15\%$) performance improvement - however we understand that this might not apply to every system.

Illustration of the remaining nine patterns are given below.

### A.3 Pattern 3

Accessing global variables or built-in functions is slightly more time consuming than local.

```
def slightly_slower(asequence, adict):
    for x in asequence:
        adict[x] = hex(x)

def slightly_faster(asequence, adict):
    myhex = hex
    for x in asequence:
        adict[x] = myhex(x)
```

The latter is faster, because it is not accessing the built-in function globally in the second loop.

### A.4 Pattern 4

Addition is faster than multiplication.

```
x+x
```

is faster than

```
x*2
```

### A.5 Pattern 5

`if` statements are expensive in Python. It is recommended to use `if` statements only if you know that they will be executed only once. If they are going to be used more than one time (e.g., if statement within a for loop) it may be efficient to replace them (if possible).

### A.6 Pattern 6

Loop unrolling - instead of a for loop:

```
def use_for():
    for i in range(1000):
        do(i)
```

use `map()`, if possible:

```
def use_map():
    map(do, range(1000))
```

### A.7 Pattern 7

Built in functions are the way to go in all cases because they are written in `C`, which makes them faster than any other logic you might use [48]. However, there may also be differences in between them, e.g. `type()` vs. `isinstance()`, and `range()` vs. `xrange()` (which is resolved in Python 3).

### A.8 Pattern 8

Do not do:

```
for key in some_dict.keys()
```

but:

```
for key in some_dict
```

### A.9 Pattern 9

Cache a method called in a loop (this is easily applied on methods that return consistent results). Method look-ups can be expensive. So for example in:

```
for i in xrange(1000):
    myobj.compute(i)
```

you can eliminate the look-up:

```
compute = myobj.compute
for i in xrange(1000):
    compute(i)
```

### A.10 Pattern 10

Swapping (memory optimization). Avoid:

```
temp = x
x = y
y = temp
```

but do:

```
x, y = y, x
```

### A.11 Pattern 11

Try to do operations in place, instead of creating a new instance (memory optimization). Do not do:

```
sortme = sorted(sortme)
```

but instead:

```
sortme.sort()
```

## ACKNOWLEDGMENTS

The authors would like to thank the RTB development team of Addictive Mobility for the invaluable discussions, suggestions and feedback they provided during this research study. The research is supported in part by MITACS Accelerate Grant # IT04452 and NSERC Discovery Grant RGPIN-2015-06075.

## REFERENCES


[1] S. Han, Y. Dang, S. Ge, D. Zhang, and T. Xie, "Performance debugging in the large via mining millions of stack traces," in *Proc. of the 34th Int. Conf. on Softw. Eng.*, 2012, pp. 145–155.

[2] G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu, "Understanding and detecting real-world performance bugs," *ACM SIGPLAN Notices*, vol. 47, no. 6, pp. 77–88, 2012.

[3] C. Killian, K. Nagaraj, S. Pervez, R. Braud, J. W. Anderson, and R. Jhala, "Finding latent performance bugs in systems implementations," in *Proc. of the eighteenth ACM SIGSOFT Int. symposium on Foundations of Softw. Eng.*, 2010, pp. 17–26.

[4] Int. Advertising Bureau 2014 revenue report. http://www.iab.com/insights/report-global-mobile-advertising-revenue-hit-31-9-billion-e23-9-billion-in-2014/, last accessed 2016-05-14.

[5] S. Zaman, B. Adams, and A. E. Hassan, "Security versus performance bugs: a case study on firefox," in *Proc. of the 8th working Conf. on mining software repositories*, 2011, pp. 93–102.

[6] ——, "A qualitative study on performance bugs," in *Mining Software Repositories (MSR), 2012 9th IEEE Working Conf. on*, 2012, pp. 199–208.

[7] M. Fowler and M. Foemmel. Continuous integration. http://martinfowler.com/articles/continuousIntegration.html, last accessed 2016-01-07.

[8] T. Hall, S. Beecham, D. Bowes, D. Gray, and S. Counsell, "A systematic literature review on fault prediction performance in softw. eng." *Softw. Eng., IEEE Trans. on*, vol. 38, no. 6, pp. 1276–1304, 2012.

[9] N. E. Fenton and M. Neil, "A critique of software defect prediction models," *Softw. Eng., IEEE Trans. on*, vol. 25, no. 5, pp. 675–689, 1999.

[10] A. Orso and G. Rothermel, "Software testing: a research travelogue (2000–2014)," in *Proc. of the on Future of Softw. Eng.*, 2014, pp. 117–132.

[11] F. Rahman, S. Khatri, E. T. Barr, and P. Devanbu, "Comparing static bug finders and statistical prediction," in *Proc. of the 36th Int. Conf. on Softw. Eng.*, 2014, pp. 424–434.

[12] B. Caglayan, A. T. Misirli, A. B. Bener, and A. Miranskyy, "Predicting defective modules in different test phases," *Software Quality Journal*, vol. 23, no. 2, pp. 205–227, 2015.

[13] B. Caglayan, A. Tosun, A. Miranskyy, A. Bener, and N. Ruffolo, "Usage of multiple prediction models based on defect categories," in *Proc. of the 6th Int. Conf. on Predictive Models in Softw. Eng.*, 2010, pp. 1–8.

[14] S. H. Kan, *Metrics and models in software quality Eng.*, 2nd ed. Addison-Wesley, 2002.

[15] A. Vetro, "Using automatic static analysis to identify technical debt," in *Proc. of the 34th Int. Conf. on Softw. Eng.*, 2012, pp. 1613–1615.

[16] I. Molyneaux, *The Art of Application Performance Testing: From Strategy to Tools*. O'Reilly Media, Inc., 2014.

[17] E. Weyuker, T. Ostrand, and R. Bell, "Comparing the effectiveness of several modeling methods for fault prediction," *Empirical Softw. Eng.*, vol. 15, no. 3, pp. 277–295, 2010.

[18] K. C. Foo, Z. M. J. Jiang, B. Adams, A. E. Hassan, Y. Zou, and P. Flora, "An industrial case study on the automated detection of performance regressions in heterogeneous environments," in *Proc. of the 37th Int. Conf. on Softw. Eng.*, 2015, pp. 159–168.

[19] A. Nistor, P. C. Chang, C. Radoi, and S. Lu, "Caramel: Detecting and fixing performance problems that have non-intrusive fixes," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, vol. 1, May 2015, pp. 902–912.

[20] "MISRA-C:2012 Guidelines for the use of the C++ language in critical systems, 2012," MISRA - The Motor Industry Software Reliability Association, Guidelines, 2013. [Online]. Available: http://www.misra-c.com

[21] "MISRA C++ Guidelines for the use of the C++ language in critical systems," MISRA - The Motor Industry Software Reliability Association, Guidelines, 2008. [Online]. Available: http://www.misra-cpp.com

[22] M. Hashemi. (2014) 10 Myths of Enterprise Python — PayPal Engineering Blog. https://www.paypal-engineering.com/2014/12/10/10-myths-of-enterprise-python/, last accessed 2016-05-14.

[23] Quotes about Python — Python.org. https://www.python.org/about/quotes/, last accessed 2016-05-14.

[24] M. Jovic, A. Adamoli, and M. Hauswirth, "Catch me if you can: performance bug detection in the wild," in *ACM SIGPLAN Notices*, vol. 46, no. 10, 2011, pp. 155–170.

[25] B. Livshits and T. Zimmermann, "Dynamine: finding common error patterns by mining software revision histories," in *ACM SIGSOFT Softw. Eng. Notes*, vol. 30, no. 5, 2005, pp. 296–305.

[26] N. Serrano and I. Ciordia, "Bugzilla, itracker, and other bug trackers," *Software, IEEE*, vol. 22, no. 2, pp. 11–13, 2005.

[27] B. Çağlayan, A. B. Bener, and A. Miranskyy, "Factors affecting team evolution during software projects," in *Proc. of the 7th Int. Workshop on Cooperative and Human Aspects of Softw. Eng.*, 2014, pp. 16–23.

[28] Z. Li, M. Gittens, S. S. Murtaza, N. H. Madhavji, A. V. Miranskyy, D. Godwin, and E. Cialini, "Analysis of pervasive multiple-component defects in a large software system," in *Proc. IEEE Int. Conf. on Software Maintenance*, 2009, pp. 265–273.

[29] B. Caglayan, A. T. Misirli, A. Miranskyy, B. Turhan, and A. Bener, "Factors characterizing reopened issues: a case study," in *Proc. of the 8th Int. Conf. on Predictive Models in Softw. Eng.*, 2012, pp. 1–10.

[30] K. Wright, *corrgram: Plot a Correlogram*, http://CRAN.R-project.org/package=corrgram, 2015, r package version 1.7.

[31] D. Gray, D. Bowes, N. Davey, Y. Sun, and B. Christianson, "The misuse of the NASA metrics data program data sets for automated software defect prediction," in *Proc. 15th Annual Conf. on Evaluation & Assessment in Softw. Eng.*, 2011, pp. 96–103.

[32] H. He, E. Garcia *et al.*, "Learning from imbalanced data," *Knowledge and Data Eng., IEEE Trans. on*, vol. 21, no. 9, pp. 1263–1284, 2009.

[33] G. E. Batista, R. C. Prati, and M. C. Monard, "A study of the behavior of several methods for balancing machine learning training data," *ACM Sigkdd Explorations Newsletter*, vol. 6, no. 1, pp. 20–29, 2004.

[34] D. Mease, A. J. Wyner, and A. Buja, "Boosted classification trees and class probability/quantile estimation," *J. Mach. Learn. Res.*, vol. 8, pp. 409–439, May 2007.

[35] M. Tan, L. Tan, S. Dara, and C. Mayeux, "Online defect prediction for imbalanced data," in *Proc. of the 37th Int. Conf. on Softw. Eng. - Volume 2*, ser. ICSE '15. Piscataway, NJ, USA: IEEE Press, 2015, pp. 99–108.

[36] R. Moser, W. Pedrycz, and G. Succi, "A comparative analysis of the efficiency of change metrics and static code attributes for defect prediction," in *Softw. Eng., 2008. ICSE'08. ACM/IEEE 30th Int. Conf. on*. IEEE, 2008, pp. 181–190.

[37] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The WEKA Data Mining Software: An Update," *SIGKDD Explor. Newsl.*, vol. 11, no. 1, pp. 10–18, 2009.

[38] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2013. [Online]. Available: http://www.R-project.org/

[39] I. H. Witten and E. Frank, *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, 2005.

[40] A. T. Mısırlı, B. Çağlayan, A. V. Miranskyy, A. Bener, and N. Ruffolo, "Different strokes for different folks: a case study on software metrics for different defect categories," in *Proc. of the 2nd Int. Workshop on Emerging Trends in Software Metrics*, 2011, pp. 45–51.

[41] M. M. Lehman, "On understanding laws, evolution, and conservation in the large-program life cycle," *J. Syst. Softw.*, vol. 1, pp. 213–221, 1980.

[42] T. Fukushima, Y. Kamei, S. McIntosh, K. Yamashita, and N. Ubayashi, "An Empirical Study of Just-in-time Defect Prediction Using Cross-project Models," in *Proceedings of the 11th Working Conference on Mining Software Repositories (MSR'14)*, 2014, pp. 172–181.

[43] A. P. Bradley, "The use of the area under the roc curve in the evaluation of machine learning algorithms," *Pattern recognition*, vol. 30, no. 7, pp. 1145–1159, 1997.

[44] R. K. Yin, *Case study research: Design and methods*, 5th ed. Sage publications, 2013.

[45] A. Porter, H. Siy, A. Mockus, and L. Votta, "Understanding the sources of variation in software inspections," *ACM Trans. Softw. Eng. Methodol.*, vol. 7, no. 1, pp. 41–79, Jan. 1998.

[46] C. F. Kemerer and M. C. Paulk, "The impact of design and code reviews on software quality: An empirical study based on psp data," *IEEE Transactions on Software Engineering*, vol. 35, no. 4, pp. 534–550, July 2009.

[47] Index of Python Enhancement Proposals (PEPs). https://www.python.org/dev/peps/, last accessed 2016-05-14.

[48] Python programming language built-in functions. https://docs.python.org/2/library/functions.html, last accessed 2016-07-27.